

How Rings assign address space

Step1: align incoming address to self (to some power of 2)

Step2: assign the result to self address

Step3: $\text{next_addr} = \text{self_addr} + \text{self_addr_space}$; // number of register used locally


Step4: send down next_addr


Example:

Dma needs 16 addr


Uart needs 4

Timer needs 256




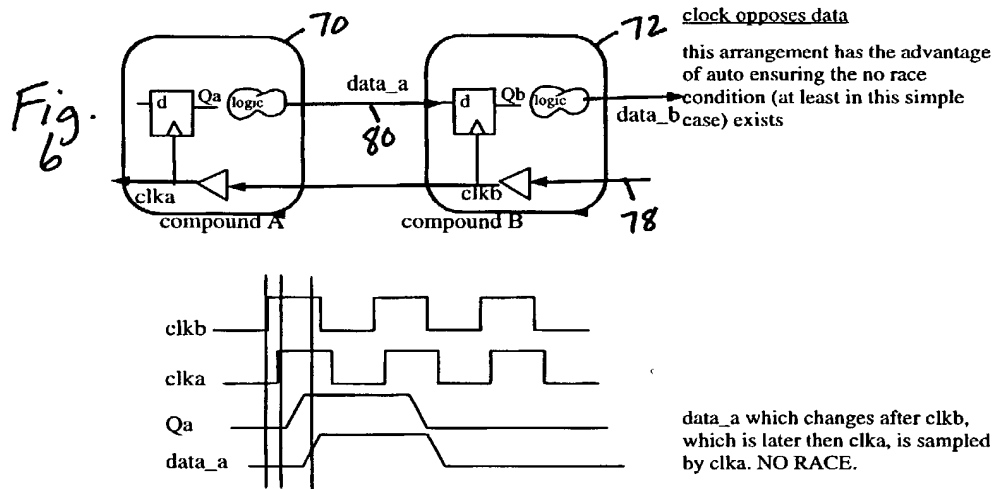


If the delay between “clk1” and “clk2” greater then the delay from Q to d of second flipflop, we have a race on our meaning right hand flipflop will sample the data of Q a whole clock period early.



clock runs with data

the problem is possible race. However, we control the logic on each flipflop leaving the compound, because it is always the same standard ring-interface module. we can ensure, that the delay will be at least enough. And more importantly easily checked after layout.





Fig.
8

Fig. 9






Fig. 12



if module "a" sends a message to module "b", ring works fine. However if most of the traffic is from "c" to "b", this is more expensive in terms of latency.

Another problem is "peak latency". Suppose that, "a" transmits mostly to "d" and "b" mostly to "c". In this case communication between "b" and "c" suffers degradation in case that peak traffic coincide.






Fig. 13




Land bridge gets its name from the fact that it is a luxury. It spans across connected modules. The idea is simple. When V2 sends message to D1 it gets to one side of the bridge. This side analyzes the destination address and by some magic (explained later) decides to short-cut the path. The message re-appears at the other end of the bridge and gets fast to D1. By same magic, message from V1 to D2 get bypassed also. message from V1 to D1 is treated directly.



Enumeration is started by "Anchor" which assigns address=1 to itself. results of enumeration are labels 1 to 7. land bridge gets two addresses, as if it were not one module. there is "near" end, that got enumeration label "3", and the "far" end marked 6.

Fig. 15



msg1 and msg2 arrive at the same time.
the bridge end must make a decision
which message to forward first.

It can be shown that unwise decision can
lead to freezout, deadlock and option price
dropping to 5\$.

Therefore MSG2 gets the priority.



Fig.
17

Fig. 18



Bridge takes responsibility for strays, but only at the "far" end. During enumeration, bridge is "polarized" to have near and far end. Near is the end first struck by enumeration message.

So we have exactly one enforcer for each ring.







Fig. 19



In land bridge ring, the situation is trickier. If V2 send message to address==5. The land bridge divert at 11/far end. it will re-appear at 3 and start cycling forever.

We have to define an algorithm that will take care of all cases.

Luckily there is a way.

Land Bridge deals only with messages arriving at the far end and being diverted. It marks and monitors only those. Messages arriving at near end, keep their markings. Messages at fdar end going through, are left alone.





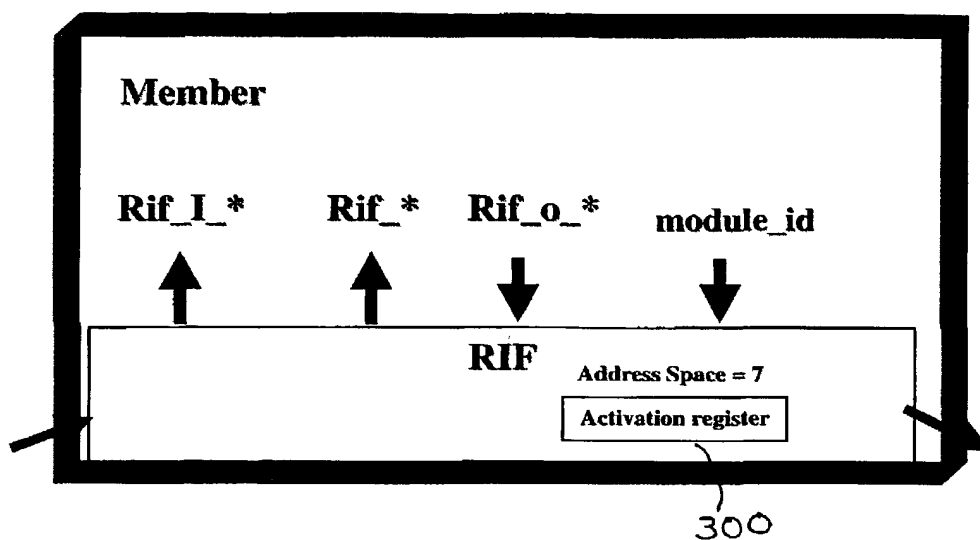


Fig. 30

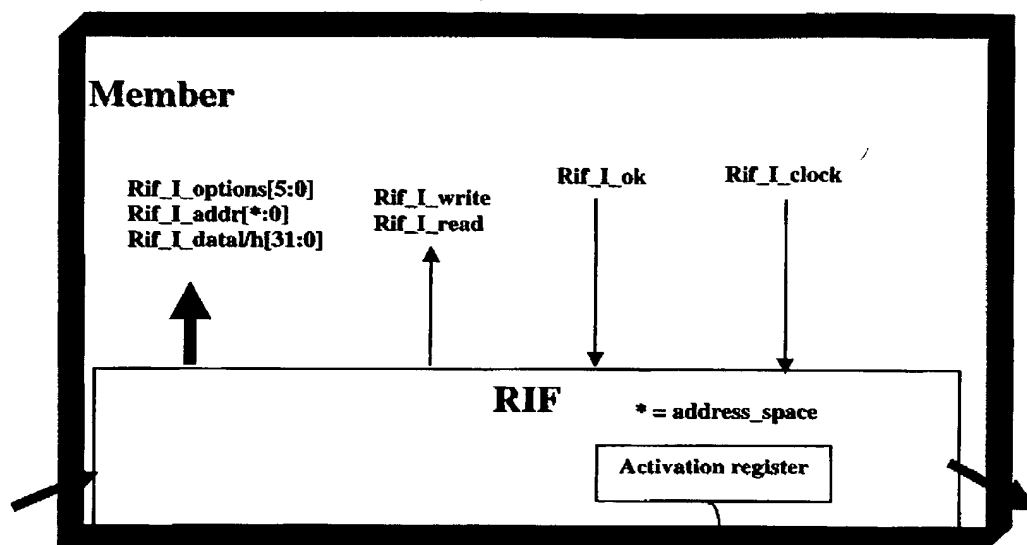


Fig. 31.

300

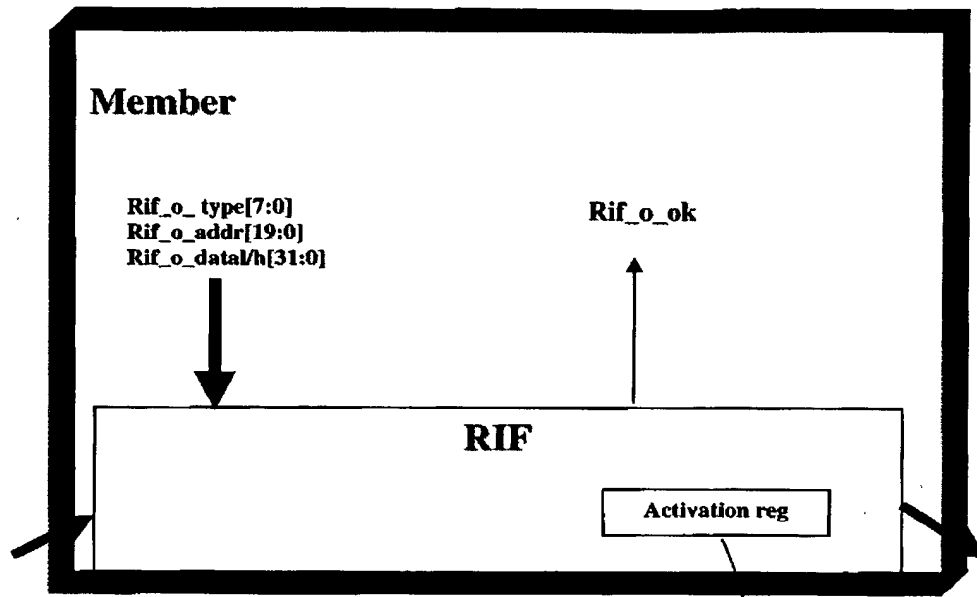


Fig. 32

300

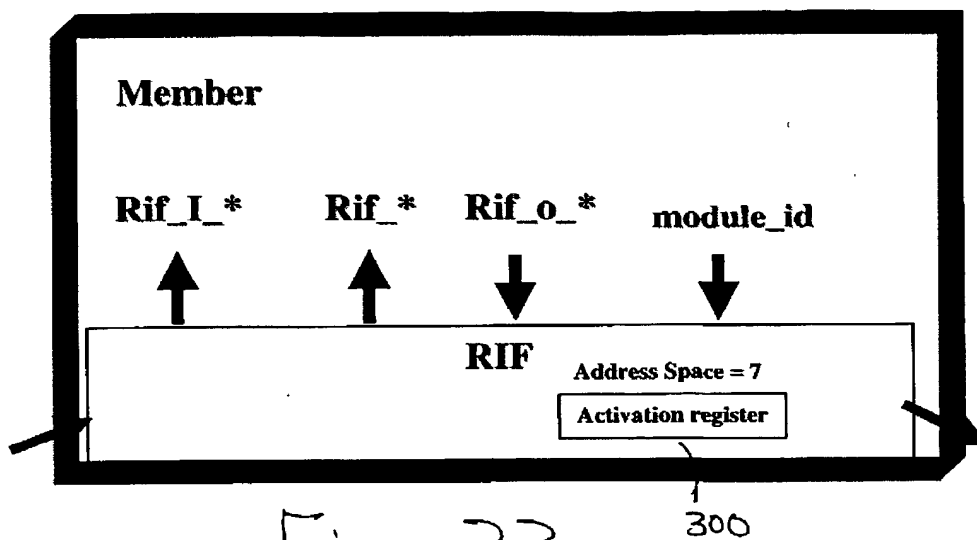




Fig. 33

Fig. 34




the second land bridge solves most traffic problems, but adds 4 clocks in the overall ring length. This is not a big problem because no message should travel the whole perimeter.

Fig. 35



The utopia interface is forced into mode that communicates in messages, not cells. We using the I/O and maybe some of the logic.

Fig. 36



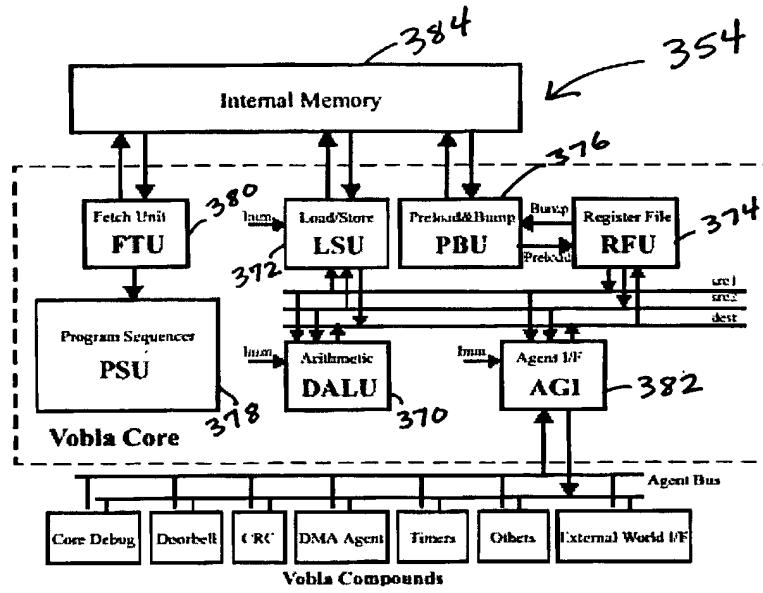



Fig.
37

Fig.
38

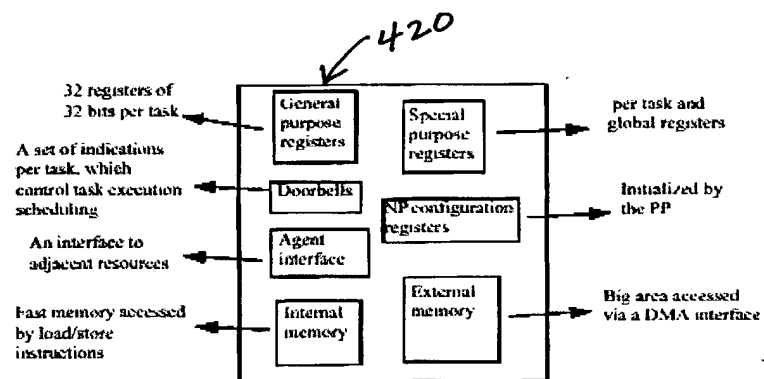
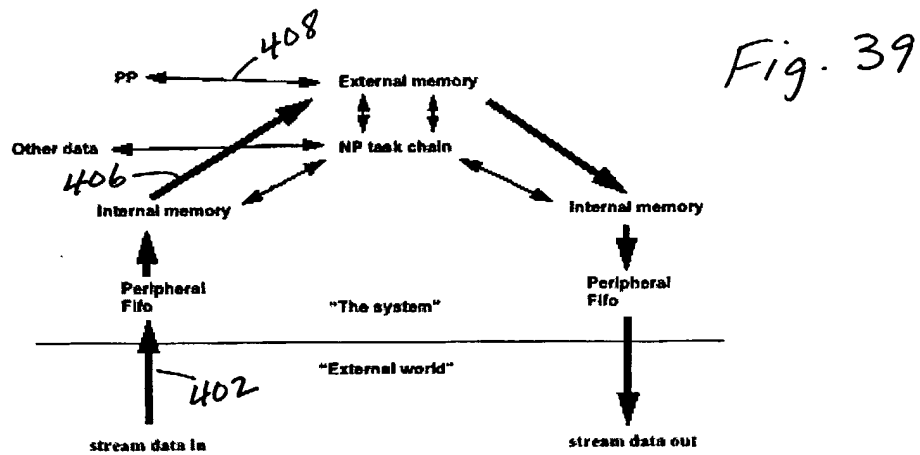


Fig. 41

R1 register:



s - sticky bit
eq - equal/zero
lt - less than/negative
gt - greater than/positive
c - carry
mb - reflection of the RAM multi-reader busy indication.

430

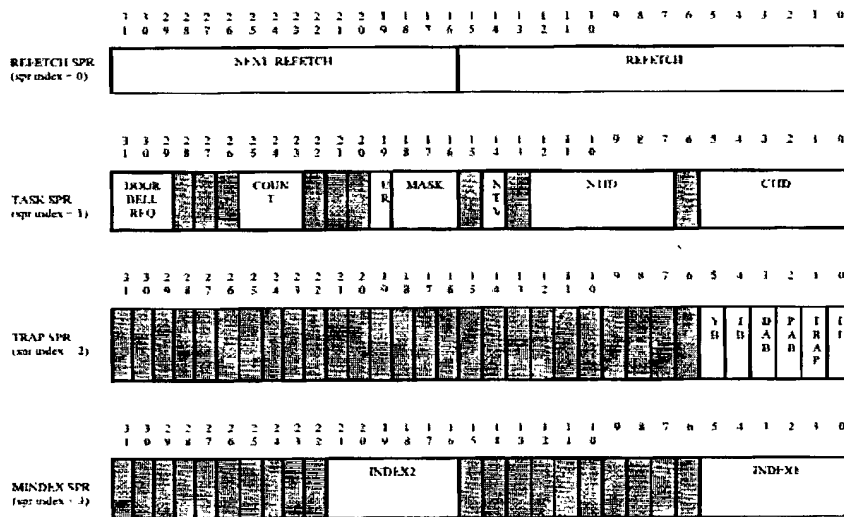


Fig. 42

440

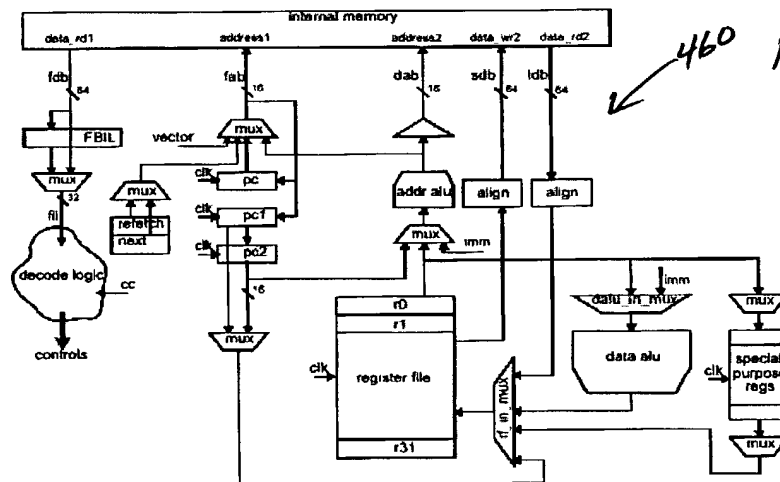
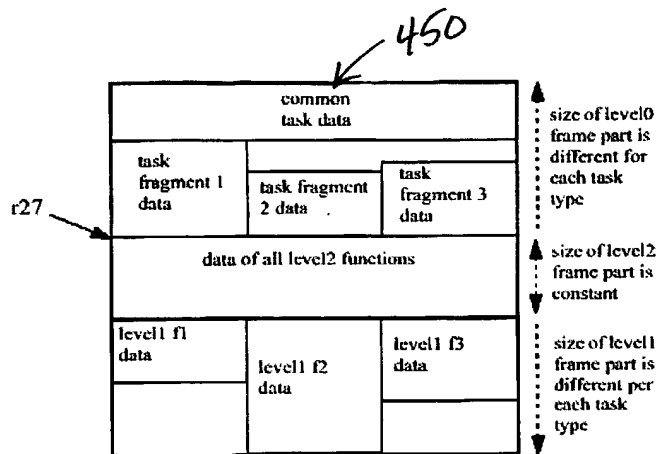
442

444

446

Frame structure
of an example
task type

Fig.
43



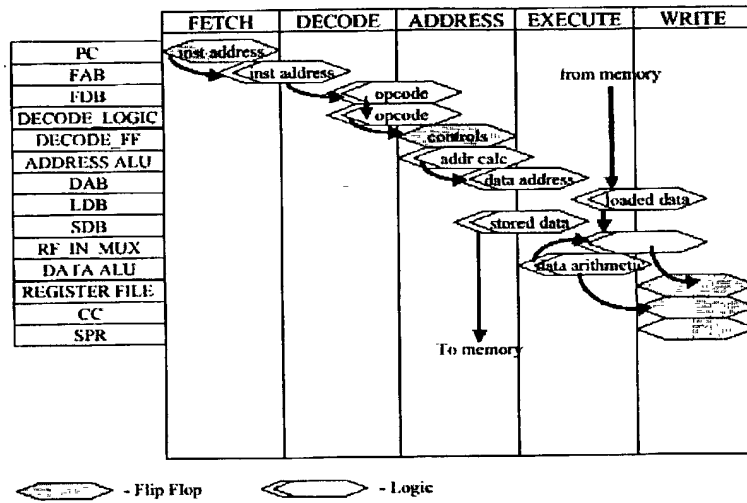
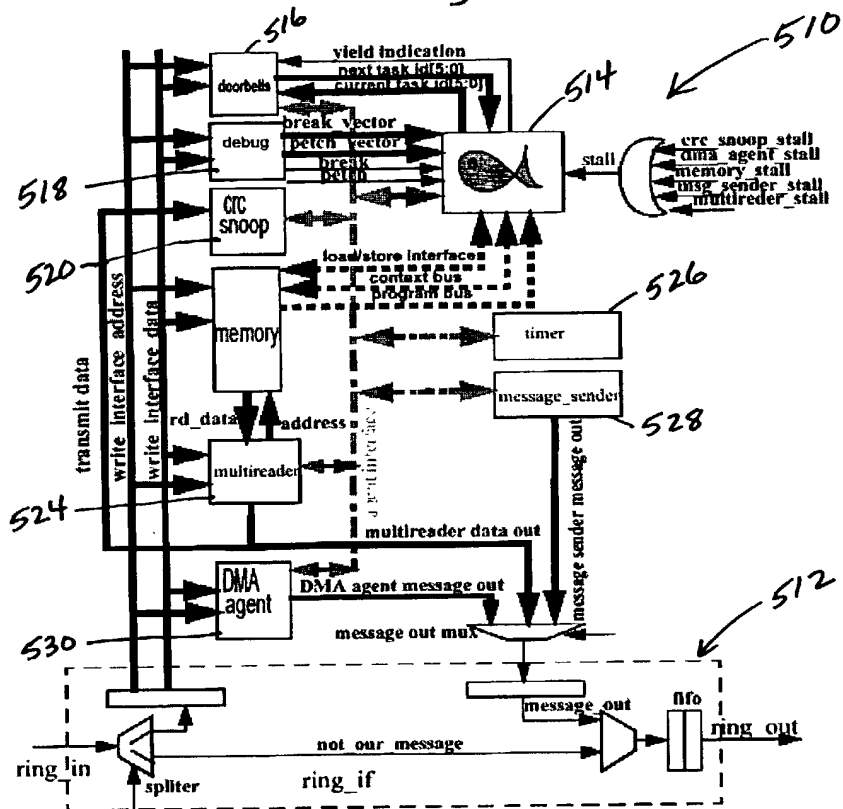


Fig. 46



Fig. 47



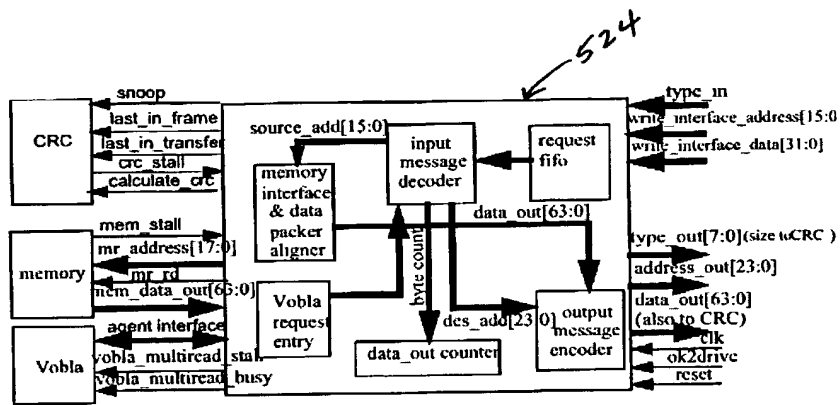


Fig.
48Fig.
49

Fig. 50

Fig. 51

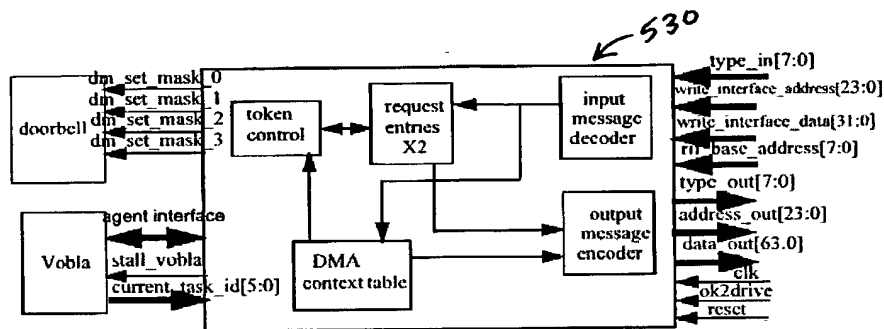
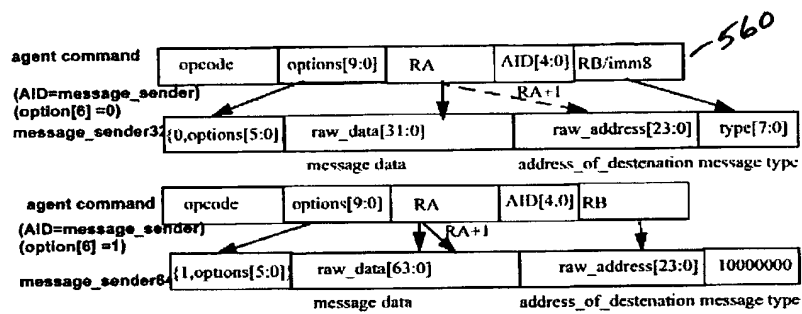
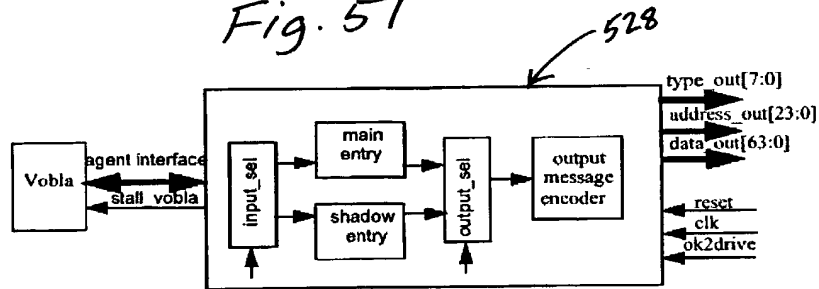


Fig. 53

Fig. 54

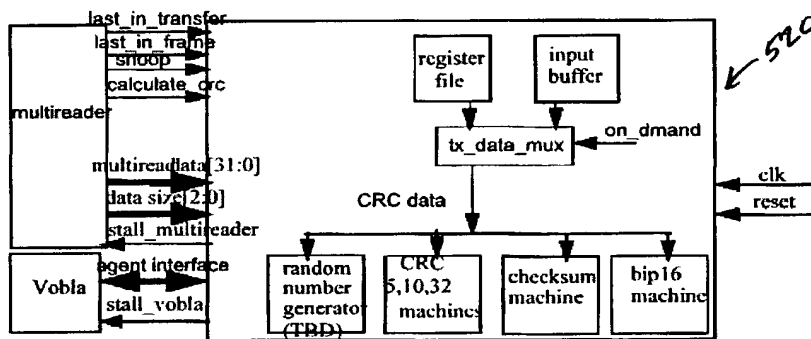


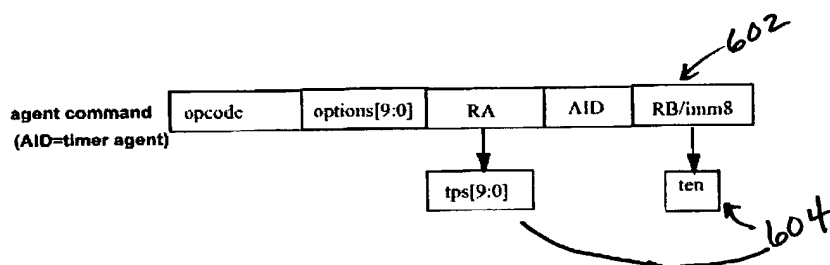
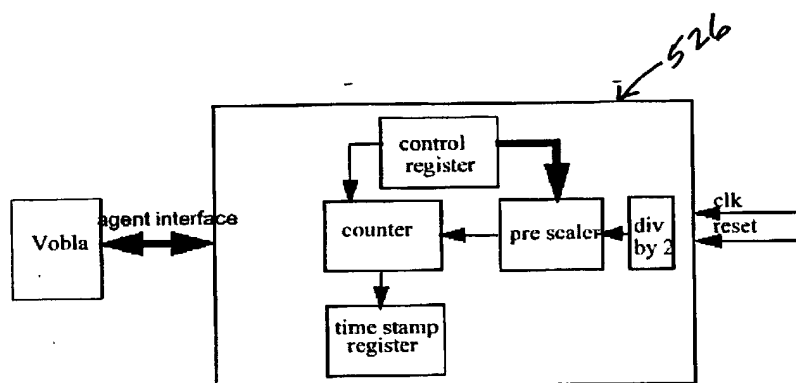


Fig. 55

Fig. 56





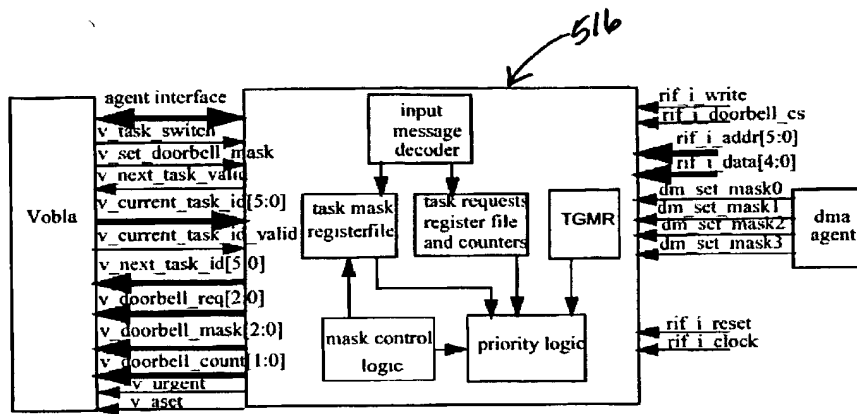


Fig. 59

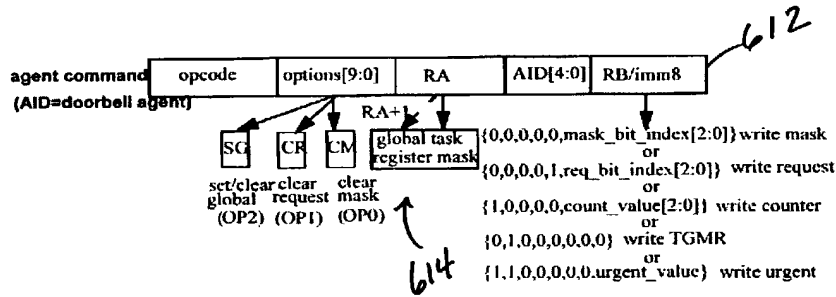


Fig. 60

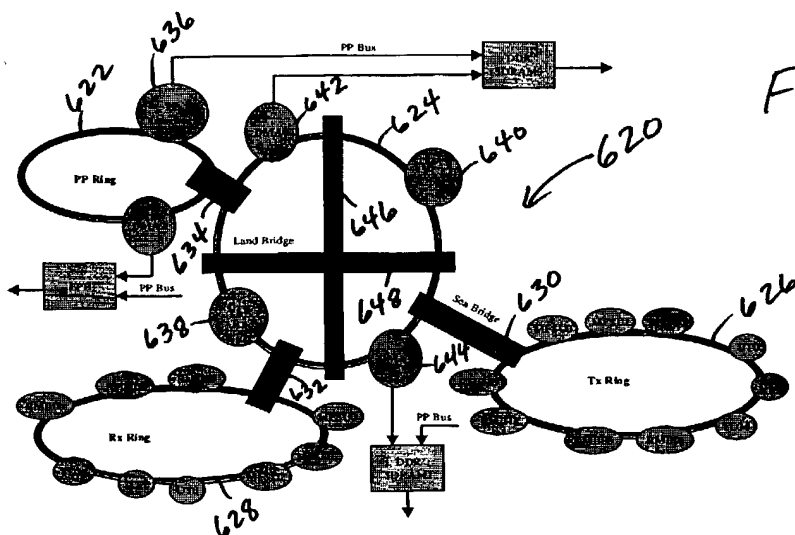


Fig. 61

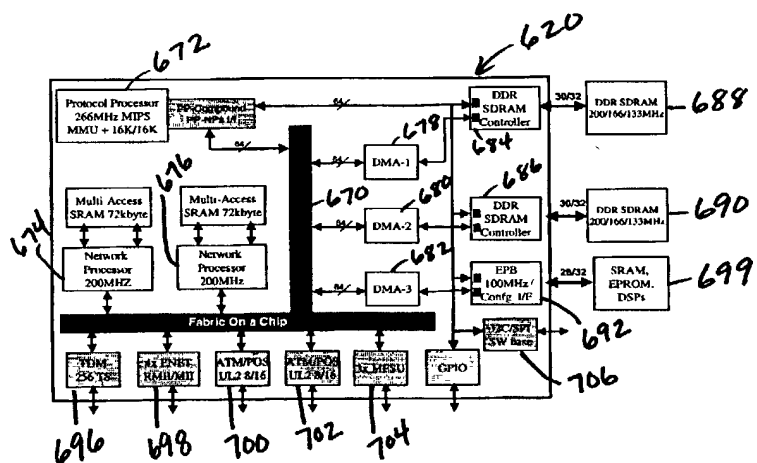



Fig.
62






Fig. 64




Fig. 65

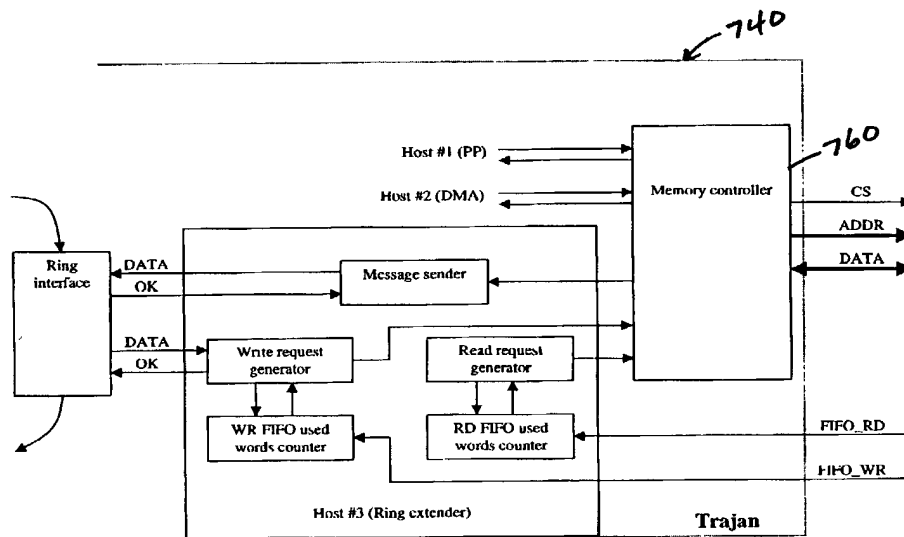

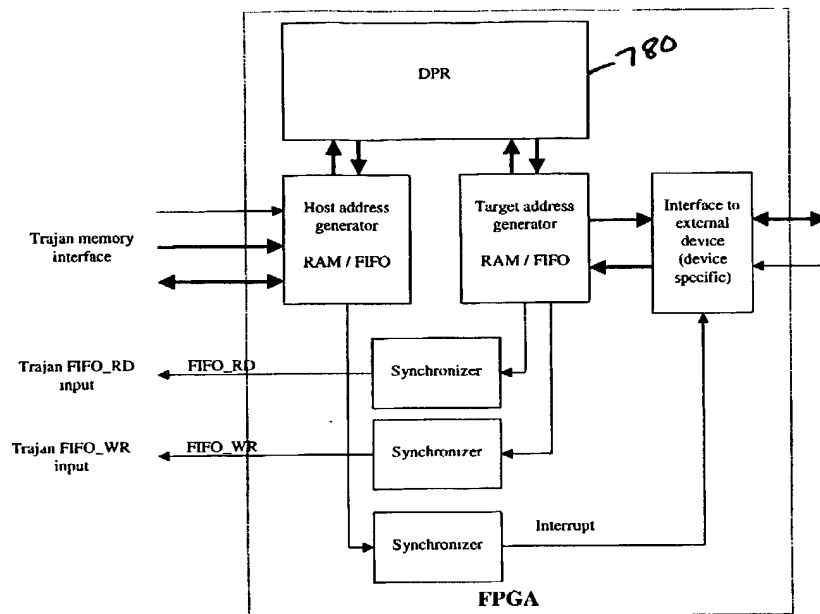


Fig. 66

*Fig. 67*

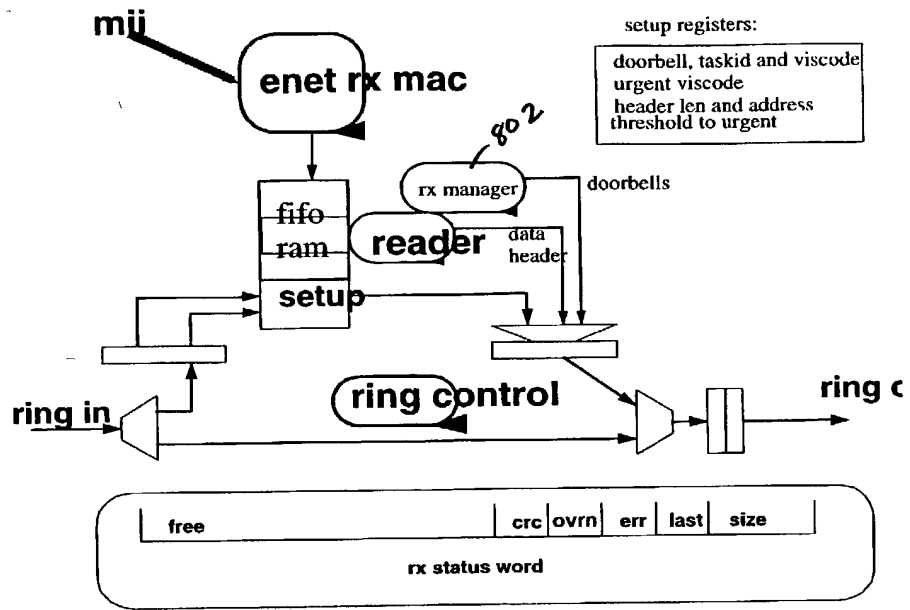


Fig. 68

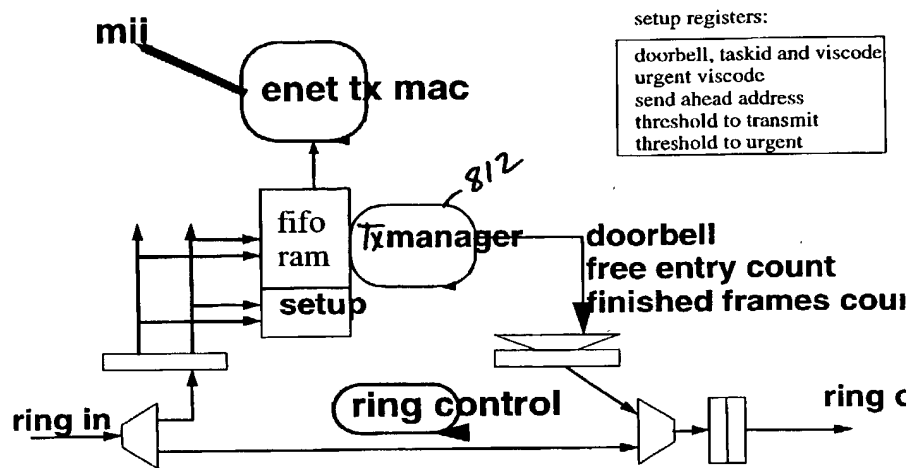


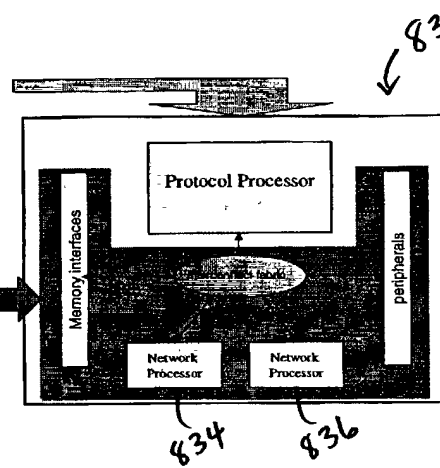
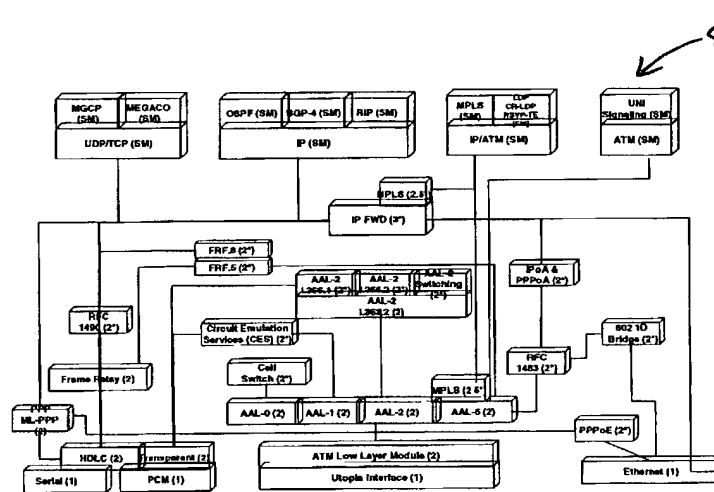
Fig. 69

Control Plane

Signaling Protocols
Protocol Management
Exception Handling
System Control &
Configuration

Data Plane

Perpacket handling
Forwarding Decision
Classification
QoS Handling
Queueing
Scheduling
Formatting

Fig.
70Fig.
71




Fig. 72

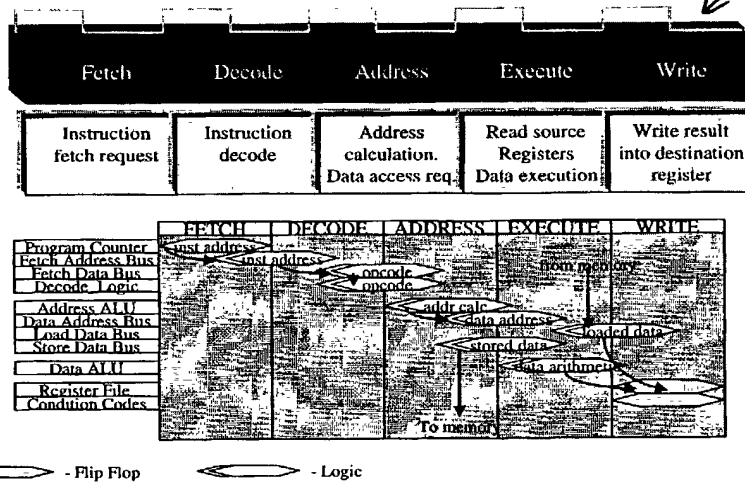
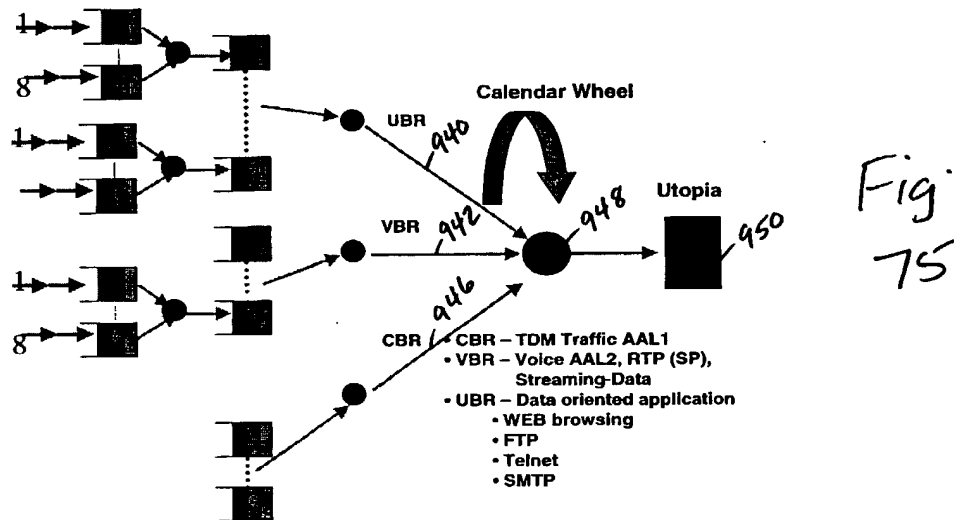



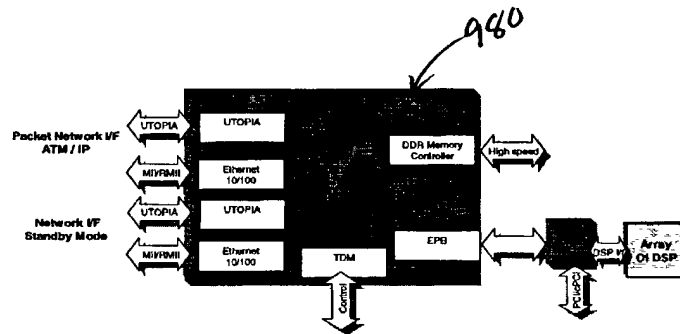
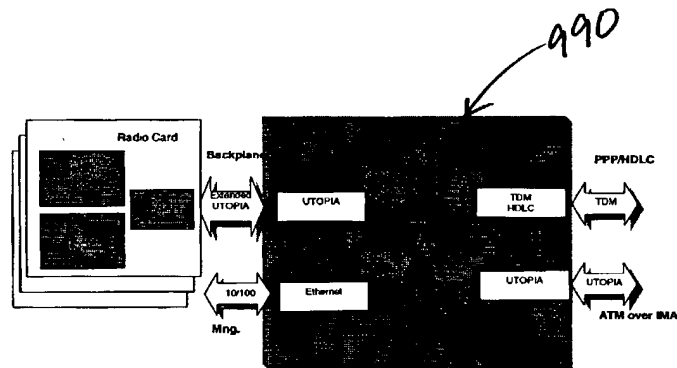


Fig. 73



Fig.
76Fig.
77

Fig.
78Fig.
79

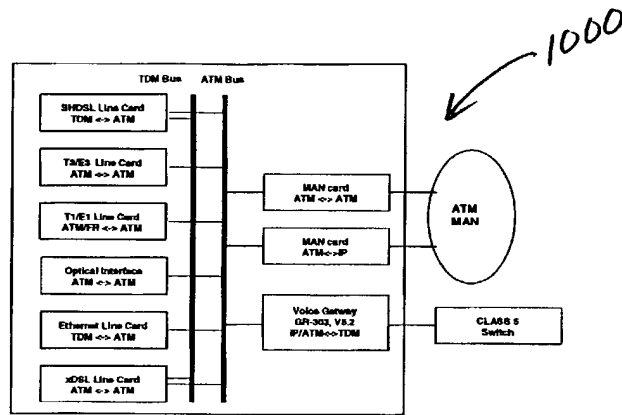


Fig. 80

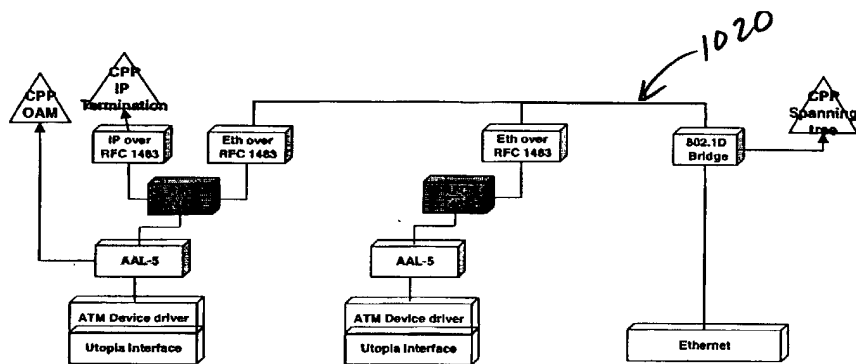


Fig. 81




Fig. 82

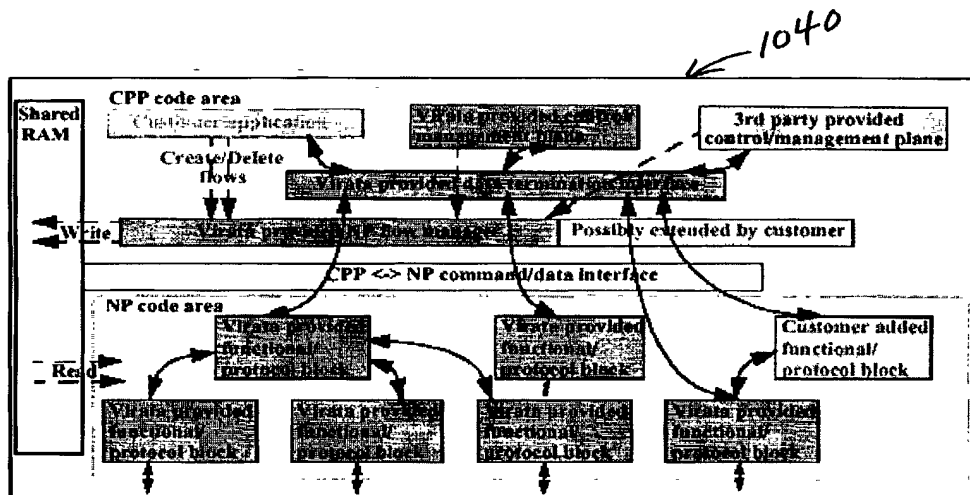


Fig. 83

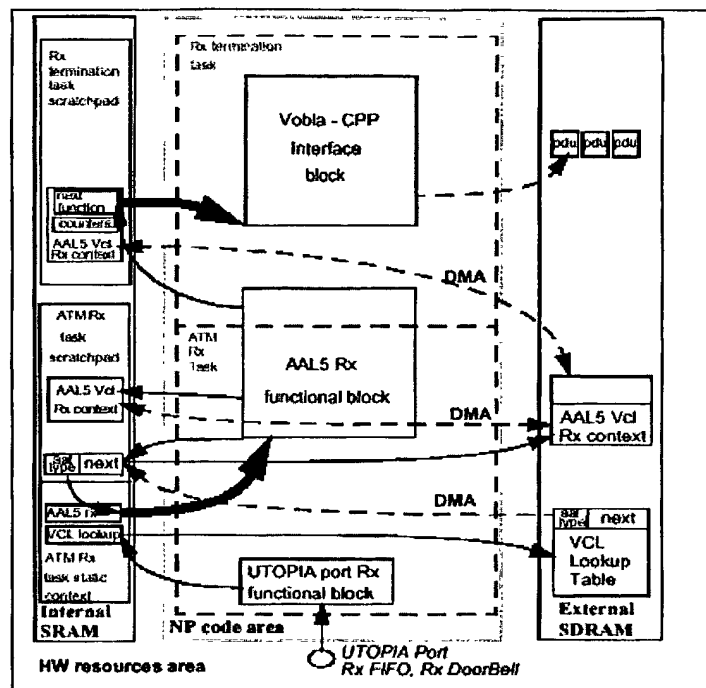
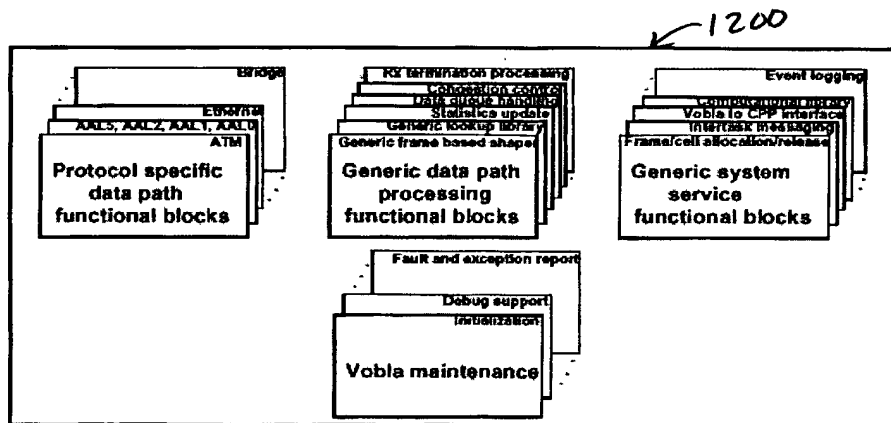



Fig. 84



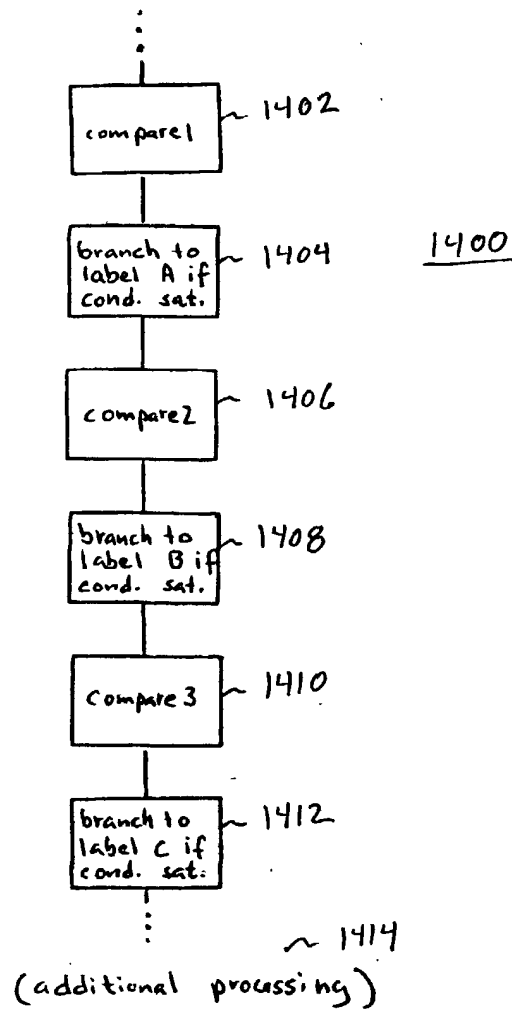


FIG. 87
(PRIOR ART)

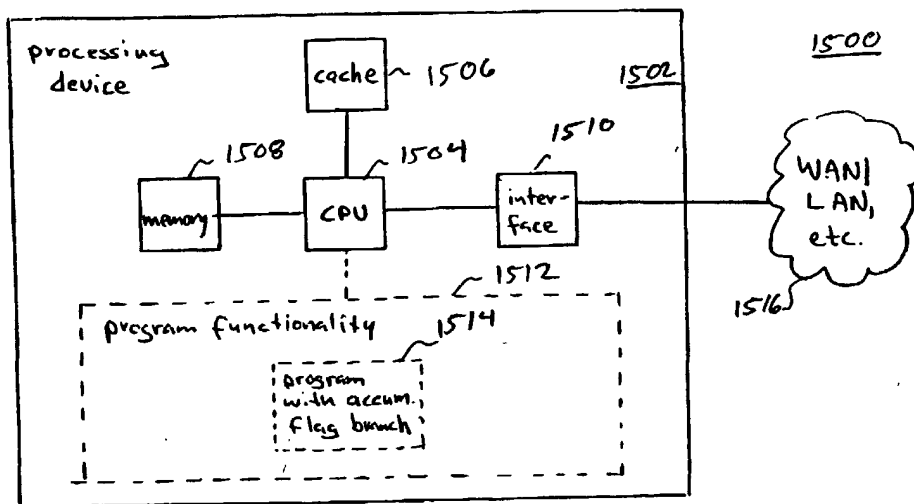


FIG. 88




FIG. 89

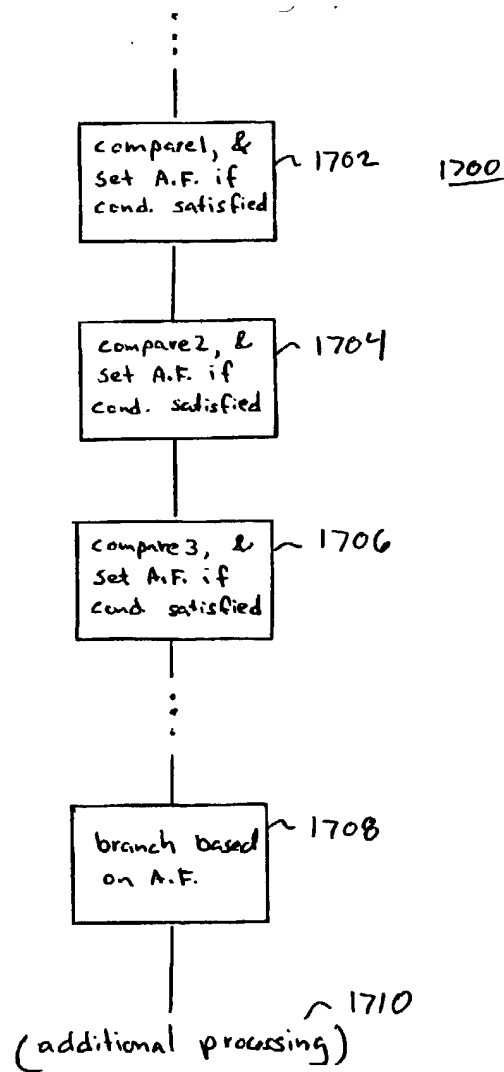


FIG. 90